

REMARKS

This application contains claims 1-26. Claims 13 and 26 have been canceled without prejudice. Claims 1-12, 14-17 and 24 are hereby amended. No new matter has been introduced. Reconsideration is respectfully requested.

Applicant thanks Examiner Tsai for the courtesy of a personal interview with Applicant's representative, Sanford T. Colb (Reg. No. 26,856), held in the USPTO on August 9, 2006. At the interview, Mr. Colb presented a proposed amendment to claim 1 and argued the patentability of the new claims over the cited art (Citron). The Examiner agreed that the proposed amendment would overcome the present rejection, subject to a detailed showing of support for all the newly-added claim elements. The Examiner suggested that claim 1 be drawn to an "integrated circuit chip," rather than an "integrated circuit device," as in the claims as filed. Applicant has amended claim 1 along the lines agreed to in the interview and has made a similar amendment to independent method claim 14.

The abstract was objected to for use of the term "m1" in place of " m_1 ". The Examiner is apparently referring to the abstract appearing in the published version of this application, US 2005/0114600. The error in the abstract was evidently a typographical error made by the USPTO, since the application as filed, as it appears in PAIR, uses " m_1 " correctly. Therefore, Applicant respectfully submits that no amendment of the abstract as filed is required.

Claims 1-26 were rejected under 35 U.S.C. 102(b) over Citron et al. ("Creating a Wider Bus Using Caching Techniques"). Applicant has amended independent claims 1 and 14, as agreed in the interview, in order to clarify the distinction of the present invention over Citron.

Dependent claims 2-12, and 15-17 have been amended for proper antecedence in view of the amendment to the independent claims (and claim 24 has been amended to correct an informality). Claims 13 and 26, whose limitations have been incorporated into claims 1 and 14, respectively, have been canceled.

Citron describes the use of caching, in tables at both ends of a bus, high-order bits of words transmitted over the bus (abstract). He is concerned explicitly with reducing "the number of pins and wires required to communicate, load, or store a digital word of information in a computer system" and solving the problem that occurs in "off-chip communication" of the mismatch between the wide internal data paths of current microprocessors and the relatively small number of pins per chip (page 90, Introduction, first two paragraphs).

To test these ideas, Citron simulated communications between a processor and a memory module via a memory bus (page 92, last three lines). In these simulations, it was assumed that the processor had an on-chip cache, because of the strong influence of the cache on external bus performance (page 93, "3.1 Simulations," second paragraph; and page 95, "3.2.2 On-chip cache present"). In one experiment, Citron compared the effect of placing bus expanders between a L1 cache and memory and between a L2 cache and memory (page 97, first column). In every case, however, the bus expansion techniques were used between chips, while data paths on each chip were assumed to be arbitrarily wide.

Embodiments of the present invention, on the other hand, are directed to bus expansion in communication between a processing component and caches on the same integrated circuit chip. This element of the invention was expressed in the preambles of the independent claims as filed, as well as in the specification. (See, for example, the abstract and paragraph 0001 in the published

version of this application, US 2005/0114600.) Applicant has amended claim 1 to include this distinction explicitly in the claim elements.

Applicant has further incorporated into claim 1 the recitation that both L1 and L2 caches are present on the chip, with respective buses and bus expanders between the processing component, the L1 cache, and the L2 cache. This configuration was recited originally in dependent claim 13 and is shown in Fig. 3 of the present patent application. It allows the widths of the address buses both between the L1 cache and the processing component and between the L1 and L2 caches to be made smaller than the memory address length m_1 . It also permits a still greater gain in bus efficiency by enabling the bus expanders in the L1 cache to share a common table for address field compaction, as described in paragraph 0045 of the specification and illustrated by cache line 106 in Fig. 3. Claim 1 has been amended to recite this table-sharing feature, as well. (In some implementations of this feature, both address buses will have the same width n_1 , as recited in the draft claim proposed at the interview; but this limitation is not essential and has been omitted from the present amendment.)

The following chart shows the correspondence between the amended language added to claim 1 and the disclosure provided in the application as filed:

Claim 1	US 2005/0114600
An integrated circuit <u>chip</u>	[0001] "buses used for data transfer in a processor chip." [0007] "In some embodiments of the present invention, bus expanders are used to reduce the width of buses that connect processing components of a

Claim 1	US 2005/0114600
	microprocessor to a cache on the microprocessor chip."
<u>a Level 1 (L1) cache and a Level 2 (L2) cache on the chip</u>	Claim 13 as filed. Fig. 3 shows L1 and L2 caches on the chip of microprocessor 70. The on-chip elements are surrounded by the dashed line in the figure (in contrast to "off-chip memory 72," noted in [0041]).
<u>first and second buses on the chip... and third and fourth buses on the chip coupled between the L1 cache and the L2 cache</u>	Claim 13 as filed. Fig. 3 shows address and data buses 94 and 96 ("first and second buses") coupled between the L1 data cache and LSU 82 (the load/store unit - a "processing component"), and additional address and data buses 94 and 96 ("third and fourth buses") coupled between the L1 data cache and the L2 cache.
<u>the first and third buses having bus widths smaller than m_1</u>	The width characteristic of the first bus was recited in claim 1 as filed. [0043] "Comparable buses connect L2 cache 88 with L1 caches 84 and 86, and the data on these buses may likewise be compacted and de-compacted by bus expanders 102 and 104." Therefore, the "third bus" will also have a width smaller than the memory address length m_1 .
<u>the L1 and L2 caches</u>	Claim 13 as filed.

Claim 1	US 2005/0114600
<u>comprising respective address bus expanders coupled to the first and third buses</u>	Fig. 3 shows bus expanders 98, 104 and 102 in the L1 and L2 caches, coupled to the "first and third buses." These bus expanders have similarly functionality, as explained in [0043].
<u>the L1 cache comprises a table for use in compaction of address fields that is shared by the address bus expanders coupled to the first and third buses</u>	[0045] "...bus expanders 98 and 104... may be configured to use the same respective tables for compaction of the appropriate address... fields." As shown in Fig. 3, these bus expanders in the L1 cache respectively serve the "first bus" (L1 cache-LSU address bus) and the "third bus" (L1 cache-L2 cache address bus), as identified above.

The combination of elements recited in claim 1, in accordance with the present amendment, is neither taught nor suggested by the cited art. Citron describes off-chip bus expansion techniques for purposes of solving problems (data width mismatch) that do not occur on-chip. Furthermore, there is no suggestion in the prior art that would have led a person of ordinary skill to the high level of integration of bus expansion elements - including sharing a common table in the L1 cache - that is recited in the amended claim.

Thus, claim 1, as amended, is believed to be patentable over the cited art. In view of the patentability of claim 1, dependent claims 2-12 are also believed to be patentable.

Independent claims 14, as amended, recites a method for coupling a processing component on an integrated

circuit chip to a Level 1 cache and a Level 2 cache on the chip. The method is based on principles similar to those of the chip recited in claim 1 (as amended). Therefore, claim 14 is believed to be patentable over the cited art, as are claims 15-25, which depend from claim 14.

Applicant believes the amendments and remarks presented above to be fully responsive to all of the objections and grounds of rejection raised by the Examiner. In view of these amendments and remarks, all of the claims now pending in this application are believed to be in condition for allowance. Prompt notice to this effect is requested.

Please charge any fees associated with this paper to deposit account No. 09-0468.

Respectfully submitted,

By: /Stephen C. Kaufman/
Stephen C. Kaufman
Reg. No. 29,551
Phone No. (914) 945-3197

Date: August 22, 2006

IBM Corporation
Intellectual Property Law Dept.
P. O. Box 218
Yorktown Heights, New York 10598